

NATIONAL INSTITUTES OF HEALTH
DIVISION OF RESEARCH RESOURCES
BIOTECHNOLOGY RESOURCES BRANCH

SECTION I - RESOURCE IDENTIFICATION

Report Period: Grant Number:
From 1 August 1974 to 31 July 1975 RR-00785-02

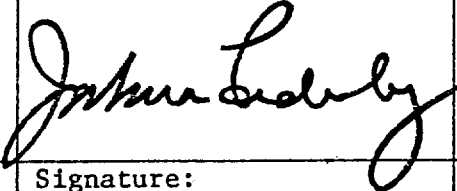
Name of Resource: Stanford University Medical Experimental Computer	Resource Address: Stanford University Stanford, California 94305	Resource Telephone Number:
Principal Investigator: Joshua Lederberg, Ph.D.	Title: Chairman and Professor Department of Genetics	Academic Department: School of Medicine Department of Genetics
Grantee Institution: Stanford University	Type of Institution: Private University	Investigator's Telephone No.: (415) 497-5801

Name of Institution's Biotechnology Resource Advisory Committee:

SUMEX-AIM Executive Committee

Membership of Biotechnology Resource Advisory Committee:

<u>Name</u>	<u>Title</u>	<u>Department</u>	<u>Institution</u>
Saul Amarel, Ph.D.	Chairman and Professor	Computer Science	Rutgers University
Donald Lindberg, M.D.	Professor Director	Pathology Information Science Group	University of Missouri School of Medicine

Principal Investigator: Joshua Lederberg, Ph.D. Chairman and Professor	Signature: 	Date: May 27, 1975
Stanford University Official:	Signature:	Date:

II RESOURCE OPERATIONS

II.A PROGRESS

II.A.1 RESOURCE SUMMARY AND GOALS

The SUMEX (Stanford University Medical EXperimental computer) project is a computer resource funded by the Biotechnology Resources Branch of the National Institutes of Health and encompasses a dual mission: 1) the promotion of applications of artificial intelligence (AI) computer science research to biological and medical problems and 2) the demonstration of computer resource sharing within a national community of health research projects. The SUMEX resource resides administratively within the Genetics Department of the Stanford University Medical School and serves as a nucleus for a growing community of projects, both within and external to Stanford. SUMEX provides computing facilities specifically tuned to the needs of AI research and communication tools to facilitate inter- and intra-group contacts as well as trial dissemination of research products to medical users. The project also develops tools for and takes an active role in stimulating community relationships among collaborating projects and medical researchers.

User projects are separately funded and autonomous in their management and are selected for access to SUMEX on the basis of their scientific and medical merits as well as their commitment to the community goals of SUMEX (see Section II.A.3 on page 18). Currently active projects span a broad range of application areas such as clinical diagnostic consultation, molecular biochemistry, belief systems modeling, mental function modeling, and instrument data interpretation (see Section IV on page 54).

Artificial Intelligence is a branch of computer science which attempts to discern the underlying principles involved in the acquisition and utilization of knowledge in reasoning, deduction, and problem-solving activities. Two recent reviews give some perspective on the current state of AI (see Nilsson, N.J., "ARTIFICIAL INTELLIGENCE", Information Processing 74, North-Holland Pub. Co. and Feigenbaum, E.A., extracts from an informal report to ARPA-IPTO, attached as Appendix A, page 115). Each authorized project in the SUMEX community is concerned in some way with the application of these principles to medical problems. The tangible objective of this approach is the production of computer programs which, using formal and informal knowledge bases together with mechanized hypothesis formation and problem solving procedures, will be more general and effective consultative tools for the clinician and medical scientist. The exhaustive search potential of computerized hypothesis formation and knowledge base utilization, constrained where appropriate by heuristic rules or interactions with the user, has already begun to produce promising results in the areas such as chemical structure elucidation, diagnostic consultation, and mental function modeling. Needless to say, much is yet to be learned in the process of

fashioning a coherent scientific discipline out of the assemblage of personal intuitions, mathematical procedures, and emerging theoretical structure of the "analysis of analysis" and of problem solving.

Our community building role is based upon the current state of computer communications technology. While far from perfected, these new capabilities offer much needed additional freedom in defining collaborative linkages, both within a given research project and among them. Several of the active projects on SUMEX are based upon the collaboration of computer and medical scientists at geographically separate institutions; separate from each other as well as from the computer resource. Another major goal of the network experiment is to enable diverse projects to interact more directly and to facilitate selective demonstrations of available programs to physicians and medical students. Even in their current developing state, such communication facilities allow access to the rather specialized SUMEX computing environment and programs from a great many areas of the country (and to some extent in Europe) for potential new research projects and for research product dissemination and demonstration.

This past year has seen the SUMEX resource become fully operational; the initially designed hardware configuration is installed, both the ARPANET and TYMNET network connections are finally installed and working, and the menu of available user software is filling out. The resource was formally dedicated in mid-November 1974 with a one day symposium held at Stanford to describe and illustrate the objectives, capabilities, and opportunities of the resource.

Over this year the complement of projects has increased from the initial group of 5 to include 9 formal projects and a group of informal pilot efforts. Already a number of examples of the benefits of inter- and intra-group interactions have come to light. There have also been substantial efforts to introduce non-affiliated research people to a number of the programs which are far enough along in their development. The management committees which help direct the allocation and development of the resource are functioning and are actively pursuing the recruitment of additional significant projects and establishing necessary resource allocation policies.

II.A.2 TECHNICAL PROGRESS

II.A.2.a SYSTEM DEVELOPMENT AND OPERATIONS

DEC PDP-10 Hardware:

At the time of the last report, the initially designed configuration was almost complete, lacking principally the high speed swapping storage and the network connections (see "Communications" on page 9 below for a summary of network status). On an interim basis, swapping was being done from the moving head disks and remote communication was handled through several IN-WATS lines to the eastern portion of the country.

A fixed-head swapping system had been selected from Digital Development Corporation (DDC) to be interfaced to a DEC Special Systems controller because of a greater capacity and transfer speed for the same money compared to equipment offered by DEC. The DDC fixed head disks also offered the option to organize the storage by pages (1 page = 512 36-bit computer words) as is desirable for TENEX operation. The technology for the device had been demonstrated prior to ordering but DDC nevertheless experienced problems in producing the needed drive and after a 4 month delay agreed to substitute a system based on an older and more expensive technology. In order to meet the capacity specifications of the original device, two of the substitute devices were required and agreed to at no price increase. The delivery was made in two stages; an interim slow device in October 1974 and the 2 final devices in early January 1975. In a very difficult situation, DDC proved to be highly cooperative and responsible in meeting their obligations.

The performance of the system improved dramatically with the installation of the fixed head swapping device as was, of course, expected. Average overhead for the system (even in the then fairly lightly loaded state) dropped from over 30% to about 10%.

As we approached the second year and facility loading increased, we projected that two aspects of the system would become bottlenecks; these were on-line file space and swapping efficiency. These projections were based on observed file space consumption and tests of system overhead under a simulated large program environment (LISP) on the IMSSS KA-TENEX system. The details of the tests and the plan (reproduced in Appendix B) suggested that substantial efficiency improvements could be achieved by adding memory and fixed-head swapping storage. The added memory allows more runnable jobs to be in memory at a given time to use cycles otherwise lost in I/O waiting to load a runnable job. The added fixed head swapping avoids overflow to moving head devices which are much slower.

The AIM Executive Committee reviewed and accepted this proposal. We have implemented the file system and memory augmentations as being immediately needed (within the lead time of hardware delivery). A current hardware configuration diagram is shown in Figure 1 (see page 111). As the high speed swapping store has been brought up to initial design goals only as of January 1975, we have delayed the addition of more fixed head disk space pending demonstration that that is the most effective place to augment the system.

We have been observing system performance over the past few months as the user load has increased and find that the loading during prime shift frequently approaches saturation in terms of system responsiveness for interactive users. This is illustrated in the diurnal loading data (see page page 35) where it can be seen that the load average peaks typically run between 5 and 10 during prime hours(1). The reports from the individual projects (see Section

(1) The "load average" signifies the number of jobs waiting in queue to be processed at a given instant: it measures the number of people awaiting service at that moment, so that responsiveness will be

IV, particularly page 80 and page 93) in the SUMEX-AIM community verify that the loading is subjectively approaching saturation with statements expressing concern over the ability to work during prime time and to be able to have physicians use the interactive programs with enough responsiveness so that their frustration does not go so high as to discourage them from further use.

We have asked other users to gauge their experiences and those of their medical collaborators against load average measurements as well. Their additional comments, along with those in the individual project reports, are summarized in Appendix C, together with a discussion to relate these subjective assessments to objective system operation. Most users express difficulty about being too precise in their judgements but generally agree that very noticeable response degradations set in when the load average gets above about 4 or 5 and that responsiveness deteriorates increasingly (non-linearly) above that.

Of course, the loading is mitigated at non-prime time although a number of INTERLISP users (particularly students) work in the evening and night to get a less loaded machine. We are especially open to users who can effectively use the night hours (e.g., Hawaiian and Houston-UK cooperations). Also we are developing a batch capability to off-load some of the daytime work which is not interactive-critical. However, the very nature of interactive computing in consultative programs is that human beings are involved and the work commitments of professional people such as physicians, imply that their main load in using the machine will be during prime time. As the user community grows over the next year or so, these problems in achieving acceptable program interaction response will become all the more acute.

Our performance evaluation efforts have been to better identify the bottlenecks to system throughput and, within the council-approved funding levels, to judiciously implement augmentations which maximize efficient use of the resources for the community. Having implemented previously proposed memory augmentations, we are now observing quite low overhead times, generally from 10 - 20%. Based on recent observations of drum space usage (see page 42 showing a recent diurnal drum space loading plot) we find under heavy load, where subjective response time is unsatisfactory per user comment, that we are just at the point of exceeding the existing capacity. However, because users are not always diligent about freeing up allocated pages when idle (by RESETting), we may be able to make more effective use of the available space by system software to transfer dormant pages to the larger moving head disk(2). We have made some estimates of

(approximately) inversely related to the load average. Two, three, or even four times as many users may be connected to the system at such times; but users typically take time out to ponder what the computer has reported, or the jobs may be preoccupied with input or output rather than the CPU.

(2) This is a particularly striking example of the trade-off between hardware and software investment. The economy of a software solution is enhanced by the ease with which such system programs can be shared with other facilities.

effectiveness of drum utilization and find that a substantial number of pages (exact amount varies widely but typical estimates may be in excess of 20-30%) are really dormant and could be moved to moving head disk without degradation. This would free up these pages on the fixed head devices allowing more effective use.

We are then faced with the fact that SUMEX is becoming response-time limited during prime hours. An analysis of existing data and user comments given in Appendix C points to two aspects of the machine configuration contributing to the bottleneck; CPU capacity and memory. These resources are closely inter-coupled in the performance of a time-sharing system as pointed out in the appendix and must be balanced in a well tuned system. We are at a reasonable balance point for the present configuration but have run out of inherent capacity to support current and anticipated peak loads. The system operates efficiently at between 15 and 20% overhead but does not have the speed to complete pending jobs fast enough to ensure adequate interactive response time. Our judgement, based on the arguments in Appendix C, is that the highest priority augmentation at this time should be in CPU capacity to alleviate this problem.

We are actively investigating ways in which CPU capacity can be augmented to eliminate this bottleneck. A preliminary plan is given in the budget explanation for the next grant year which proposes upgrading the present KI-10 CPU to a KL-10 (See page 51). At this time, before KL-10 deliveries have even begun, there are a number of technical uncertainties in the plan which we are working to resolve. We feel it is very important to maintain a degree of flexibility in being able to respond to needed augmentations to eliminate such bottlenecks as the community grows and its needs become more clearly defined. Consistent with this view, we request that unobligated money from year 02 be carried forward so that when the CPU augmentation plan is refined and reviewed by the Executive Committee and BRB, this money may be used to provide the needed additional processing capability for the SUMEX-AIM community.

TENEX Software - Monitor:

SUMEX is running release 1.31 of the TENEX system with modifications to accommodate the KI-10 paging hardware. Paging on the KI-10 was introduced by DEC after BB&N's (Bolt, Beranek, and Newman) experience with PDP-10 paging using a BB&N-designed pager on the older KA-10. Unfortunately, DEC did not incorporate all of the page handling hardware features of the BB&N device which facilitate demand paging. As a result, some of the hardware features for dynamically determining which pages have been modified, for changing user context, and for specifying per page access status are missing and must be simulated in software. Whereas the KI-10 hardware is intrinsically about 2 times faster than the KA-10, this additional overhead reduces the effective speed ratio to between 1.5 and 1.8 depending on the load (under light load the higher figures are achievable). Over the past year, Mr. Rainer Schulz has made important improvements to the KI-TENEX paging software and scheduler logic to significantly reduce the software overhead under heavier loads thereby better approaching the higher speed ratio.

In addition to the KI-TENEX performance improvements, we have also written new sections of code to interface the high speed swapping devices, to accommodate the network interfaces (see "Communications" on page 9 below), to control to a first approximation the CPU allocation to users, and to allow alphameric account specifications for a more transparent scheme to account for facility use among the various projects and communities. The swapping storage handler offers several features including management of multiple devices, full use of the hardware command queuing features of the DEC special systems controller, and on-line diagnostic exercising.

The resource allocation scheme includes at this time primitive facilities to control the amount of CPU time an individual user can receive consistent with the system load. Each user is assigned a percentage which defines the absolute fraction of the machine he is nominally entitled to. As he exceeds this amount, his priority within the system is penalized more and more. The time scale for this adjustment is currently 90 seconds - after each such period, penalties are reset and competition begins anew. As a given user's priority decreases, other jobs will be run preferentially if possible. If there are no other runnable jobs, then the available time is allocated to users over their aliquot so as not to waste machine capacity.

We are continuing to investigate the appropriate policies for allocation control. In particular the use of absolute versus relative priorities implies that if no one achieves his allotment, then competition essentially reverts to a laissez faire system. Also, it is becoming clear that people with fixed personal schedules (program demonstrations or busy clinicians) require some sort of reservation capability so that they can use the machine with reasonable responsiveness when they can arrange time rather than when the machine is relatively lightly loaded. We will continue to investigate the best approach to this problem in conjunction with the policy views of the management committees.

There, of course, have been periodic bugs in the software (reliability data are detailed below under "Reliability", see page 12) occasioned for the most part by monitor development efforts (network installation, swapping storage installation, etc.) but also present to a much less degree in the basic code. Aside from design oversights in development activities, the other bugs have been caused primarily by incomplete argument checking within system functions and calls - the intended use of a given section of code is often extended by the energetic user revealing these problems. It is a truism that the best way to check a system out beyond the basic functional state is to let time-sharing users on - on the whole, TENEX has borne up very well under such stress.

Operationally we have put a number of aids in the system to assure less operator error in bringing the system up. These include making sure the communications interfaces are enabled, automatically setting the time of day from other machines on the ARPANET when possible, automatically setting preplanned system halts to give users a maximum warning, various device exercisers with human readable error logging and decoding, and continuous system load monitoring and recording for diagnostic and management information purposes.

TENEX Software - Executive:

Another area of software development is in the Executive program which is the basic user interface to manipulate files, directories, and devices; control job and terminal parameter settings; observe job and system status; and execute public and private programs. As with all system work, we face a dilemma which is particularly strongly felt in this area; should we run a "standard" system or should we adapt things to user community needs and thereby tend to a "home-brew" system? This is a difficult issue in that in many respects the SUMEX community is special - it includes a broad spectrum of users from professional computer scientists and programmers to biomedical research scientists and clinicians. The latter group, of course, want a minimum impedance to using the performance programs they are interested in while the former group wants a rich assortment of system facilities and as much flexibility as possible. Since most systems are designed for the programmer community, we have adopted the viewpoint that controlled augmentations of the system must be made to accommodate the medical user. Much of this work is still in process and will be for some time. The key point of this effort is to introduce knowledge about the individual user into the system (such as his usual defaults in using system functions, his level of expertise coupled to on-line assistance, his domain of interest to alert him to new information and perhaps personalized system commands or macros convenient to his needs) so that he perceives a system tailored to his style and conventions in using the computer.

Within the existing staff we have made only initial progress toward defining our goals and implementation. We have a proposal pending with ARPA (in conjunction with members of the Computer Science Department and the Institute for Mathematical Studies in the Social Sciences) to augment the staff to concentrate more effort on this crucial interface. The name adopted is "intelligent terminal" project; in the long term with micro-computers coming of age, it is likely that a considerable amount of such individual user adaptation will reside in the user's terminal. This off-loads much overhead from central computing resources and places at the user's disposal uniform access to the range of resources tied together by networks - the "intelligent" terminal could have all of the detailed information about linking to various facilities available and ease the user's need to remember a variety of different protocols. At the same time it provides relatively uniform access to these resources, routine clerical tasks such as mail manipulation, calendar management and text editing could be handled.

We will continue to devote effort in this area in up-coming work. To date, we have made a few strides to allow the system to remember user subcommand selections through a session (such as in SYSTAT and LIST), to offer a user his own specified sequence of operations to be performed upon login (system status, reminders for today's activities, etc. - commands are stored in his own LOGIN.CMD file), and to add additional commands to the EXEC for user convenience. These include easy password modification and file preservation control (PURGE to allow combined DELETE and EXPUNGE on specific files and RETAIN to allow control over file version retention

and excess version disposition). Hardlined terminals are now automatically recognized so users do not have to specify type and this capability will be added where possible to network logins as well. The EDIT command has been modified to allow user selection of a preferred editor program (TECO, SOS, or TV) and to remember the filename and editor for future calls. We have also made the various status commands (JOBSTAT, USESTAT, and SYSTAT) more informative.

Another aspect of the EXEC we are working on is that of security and access control. We have diverse needs; regular users with valid access to all facility capabilities, guest users (principally physicians and scientists) who want to try out various performance programs applicable to their field, and other guests (primarily from other network facilities) interested in our system and software they may want to obtain. Within the file access and descriptor blocks of TENEX, we are setting up several classes of user: authorized users, guests, and network visitors. The guest facility is a simplified login procedure not requiring a name previously given to the system (we request name and affiliation data for our records and to ease future access) but requiring a password obtained from an authorized system or collaborator project staff member. The guest login will have access to a limited domain of programs - primarily message communication programs and working AI systems (e.g., chemical structure generator, MYCIN, glaucoma programs, PARRY, etc.). Guests will not be able to do general text editing, file manipulation, or program development.

The network visitor will have access only to specified files which are ready and approved for export. This implements our obligation to keep licensed software from being exported without vendor approval and at the same time offers reciprocity in software exchange which is a mainstay of the network community.

It should be pointed out that security systems are really only effective against relatively benevolent users. Many of the security schemes depend on the combinatorics of guessing passwords, and by writing clever and persistent programs can be circumvented. We have done what seems reasonable to prevent such occurrences both prospectively and by looking out for unusual activities in real time and retrospectively (The system now gives each user the most recent previous login time to help him spot possibly unauthorized use).

Communications:

A most crucial aspect of the SUMEX system is effective communication with remote users. From the user's viewpoint, the reality of using a remote computer as if it were next door depends almost singly on the ability to achieve the subjective feeling that a network connection is like a local telephone call to the computer. One way of achieving this goal is, of course, to hook up individual telephone lines for users at various places around the country. For the complement and geographical distribution of users contemplated in the SUMEX community, this would be prohibitively expensive (somewhere in excess of \$10,000 - \$15,000 per month based on early loading

expectations and sure to grow with time). In addition to these economic arguments for terminal access, networking offers other advantages for shared computing such as uniform user access to multiple machines and special purpose resources, convenient file transfers for software sharing and multiple machine use, more effective backup, co-processing between remote machines, and improved inter-user communications. We have therefore based our remote communication services on the existing networks - TYMNET and ARPANET - which allow foreign host access. These two networks complement each other; the TYMNET providing primarily terminal service with very broad geographical coverage and the ARPANET having more limited access but providing a broader range of communication services. Together, these networks give a good view of the current strengths and weaknesses of this approach.

Most of the experience to date has focussed on user terminal access to SUMEX-AIM via networks - primarily TYMNET. Our recent connection to the ARPANET has not been operational long enough to extensively assess its performance for this report although new cooperative relationships are in the process of being explored (e.g., remote file storage, processor backup, and multiple machine use). Also, for particular groups with especially convenient access to the ARPANET (Rutgers and Higher Mental Functions), work on SUMEX has been facilitated through the ARPANET connection.

Current network terminal facilities are not able to accomplish the illusion of a local call completely. Data loss is not a problem in network communications - in fact with the more extensive error checking schemes, data integrity is much higher than for a long distance phone link. On the other hand, networking has as its underlying principle that through shared community use of telephone lines, widespread geographical coverage is possible at substantially reduced cost. Our experience with individual telephone lines (IN-WATS), maintained for interim service until network facilities became operational, and network facilities bear out the cost advantages and attendant problems.

To operate 4 lines (at most 4 simultaneous users) to the east coast area cost up to \$6,000 per month including extra hour use fees. The corresponding network (TYMNET) charges are down by a factor of about 2-3 from that for a peak user load 1.5 - 2 times higher. The other side of the coin is that networks such as TYMNET are a complex interconnection of nodes and lines spanning the country (see Figure 2, page 112). The primary cause of delay in passing a message through the network is the time to transfer a message from node to node and the scheduling of this traffic over multiplexed lines. This latter effect only becomes important in heavily loaded situations; the former is always present. Clearly from the user viewpoint, the best situation is to have as few nodes as possible between him and the host - this means many interconnecting lines through the network and correspondingly higher costs for TYMSHARE, a profit-making company. Herein is the tension; to balance the unit cost of network operation against user acceptable response times. TENEX in some ways emphasizes this conflict more than other time-sharing systems because of the highly interactive nature of terminal handling (e.g., command and file

name recognition and non-printing program commands as in text editors or INTERLISP). We have connected SUMEX to the TYMNET in two places as shown in Figure 2 so as to allow more direct access from different parts of the country. Also local lines to more strategic terminal nodes are being considered for users in areas poorly served by the existing line layout.

The ARPANET, while designed for more general information transfer than purely terminal handling, has similar bottleneck problems in its topology (see Figure 3, page 113). These are reduced by the use of relatively higher speed interconnection lines (50 K baud instead of 2400 - 9600 baud lines in TYMNET) but response delays through many nodes become objectionable eventually as well. Such response problems have led to the installation of an IMP at SUMEX-AIM as related below.

We take very seriously the responsibility to provide effective communication capabilities to SUMEX-AIM users and will continuously look for ways to improve our existing facilities as well as investigate alternatives becoming available.

Technically speaking, the TYMNET connection was installed in late August 1974 and the hardware and software debugged during September. We began offering "routine" service during October and TYMNET use has increased steadily since then (see "Summary of Resource Usage", page 33). The interface is built around a Varian 620-L mini-computer supplied by TYMSHARE, Inc. with software to communicate with the other nodes in the network. Connection to the SUMEX KI-10 is through a direct memory port with an MX-10 multiplexor. The memory access improves character handling efficiency as far as the PDP-10 is concerned and allows aggregate high speed communication without excessive I/O bus loading. The TENEX software support to handle the shared ring buffers and to manage protocols between the PDP-10 and the TYMNET was developed by Mr Michael Heathman of the SUMEX project. This effort required several man-months. We have recently improved the measurement tools available to assess network responsiveness. As noted above (Figure 2), we have 2 4800 baud connections into the TYMNET to gain more direct access to the major trunks running from San Francisco to the Washington D.C. and New York areas.

We have had more than the expected share of hardware difficulties with the TYMNET interface. These have arisen primarily because the backplane wiring of the interface as supplied by TYMSHARE was too tight causing wire insulation to break around sharp bends at wire-wrap pins and causing logic element and power supply short circuits. These problems have gradually subsided as the most troublesome wires have been replaced as problems come up.

The ARPANET connection has been the subject of much administrative discussion within ARPA during the previous year and was resolved (so it appeared then) at the start of our second grant year by ARPA giving us permission to connect as a very distant host (VDH) to one of the other IMP's on the network. Whereas this was clearly suboptimal from many technical points of view, a VDH connection was much better than no connection so we proceeded to order the necessary

hardware and to design the software changes to TENEX. The VDH connection of a PDP-10 TENEX system had never been done before (other VDH's are generally small machines such as PDP-11's). This required a special interface design by BB&N and extensive interrupt level coding changes to the TENEX system to accommodate the VDH (approximately 4 man-months of work). The hardware and communication lines were installed by early January and debugging extended until late February. The time-critical hardware/software interactions required for the VDH protocols caused numerous problems in achieving a working design and produced a substantial overhead in the KI-10 when finally running. The VDH interface was finally operational in early March.

At that time, the combination of the increased load on the network IMP to which we were connected (coincidentally also located at TYMSHARE) together with the already slow response the ARPA office was experiencing in doing their computing on the OFFICE-1 computer (also on the TYMSHARE IMP), encouraged ARPA to reconsider the placement of an IMP at SUMEX. We pointed out (see Figure 3) that an IMP at SUMEX connected to the TYMSHARE IMP and also connected to the Stanford AI Lab IMP, would eliminate 4 of the 14 nodes between ARPA in Washington and the OFFICE-1 machine. ARPA agreed to this plan and supplied us with an IMP on which we have been operational as a local host since middle April. Because of the way the KI-10 interface was designed for the VDH connection (included a modified local host interface as a subelement), we were able to adapt the interface to be a local host with about 30 wiring changes and no additional cost. The change between being a VDH and operating as a local host with an IMP was dramatic. What were previously very sluggish communications, even between SUMEX-AIM and other hosts in the area (e.g., SRI and Stanford AI Lab), improved by a factor of from 3 to 5 in responsiveness and speed. We are still in the process of arranging for the installation of the additional line to the Stanford AI Lab which should be done by mid-summer.

We are being somewhat restrictive about the use of the ARPANET at the present time because of the developing policy position for the administration of the network. The administration will pass from the ARPA Information Processing Techniques Office to the Defense Communications Agency as of July 1975. At that time we expect new policies to be announced relating to access authorization and network usage cost allocation. Until these issues become clarified, we have protected the facilities for calling from SUMEX out to other sites on the ARPANET, allowing only those users who are affiliated with on-going ARPA contracts to use the facility. This also protects the SUMEX-AIM machine from acting as an expensive terminal handler for other machines - this function is better fulfilled by dedicated terminal handling machines (TIPS). All other facilities of the network connection (calling into SUMEX from anywhere on the ARPANET and FTPing files in or out of SUMEX) are available to anyone possessing an authorized directory and password for the SUMEX machine.

Reliability and Backup:

System reliability has been somewhat variable over the past

year; excellent under stable hardware and software conditions and degrading during monitor development periods (network interfacing, swapping storage installation, etc.) and during periods of hardware problems. The pertinent data are given below with indications of eras during which development took place.

SUMEX-AIM CRASH FREQUENCY (crashes/month)
AND DOWN-TIME DATA (hours/month)

Crash Type	AUG	SEP	OCT	NOV	DEC	JAN	FEB	MAR
DEC HARDWARE	14	6	7	6	17	21	13	15
SOFTWARE	2	5	0	2	1	2	8	4
ENVIRONMENT	2	3	1	0	0	0	1	1
TYMNET HDWRE	0	9	8	5	5	0	7	0
UNKNOWN	0	3	0	0	0	0	1	1
DOWN-TIME								
SCHEDULED	105	78	69	43	87	130	161	134
UNSCHED	73	39	29	19	36	21	31	31

DEFINITIONS:

Crash = Any occasion on which an operational system must be restarted or reloaded. Multiple crashes while trying to reload are not counted unless the system comes up fully between crashes.

DEC Hardware Crash = Any crash caused by a failure in the PDP-10 hardware or peripheral equipment (CPU, disk, drum, etc.)

Software Crash = Any crash caused by a malfunction within the TENEX software system.

Environmental Crash = Any crash caused by power failure, air conditioning outage, lightning, etc.

TYMNET Hardware Crash = Any crash caused by the TYMNET hardware or the interface to the PDP-10. This includes only the times when a TYMNET problem causes the PDP-10 to crash and not the times when the TYMNET goes down and the PDP-10 continues in operation.

Unknown Crash = All other crashes in which the cause is not assignable.

Scheduled Down-time = Preventive maintenance time (6-8 hours/week), scheduled maintenance to repair non-critical component failures, and system development activities requiring a stand-alone machine.

Unscheduled Down-time = Time lost because of unexpected hardware or software failure. For the most part this is the time to diagnose and either repair the problem or to reconfigure the system and bring it up to run in a somewhat degraded mode until a later scheduled shutdown for permanent repair.

DEVELOPMENT ACTIVITIES AFFECTING RELIABILITY

Whenever development efforts are undertaken which affect the monitor, some period of unreliability may result causing more crashes than are representative of the overall reliability of the system. The following gives some insight into these development efforts as reflected in the above data.

Aug -> Sept/Oct 1974: TYMNET development and installation

Early Nov 1974 and Jan 1975: Drum code development and hardware installation.

Feb/Mar 1975: Installation and checkout of ARPANET VDH code

As can be seen, we have had periods of rather serious hardware unreliability stemming from highly intermittent problems. There were a series of infant failures in the DDC swapping devices requiring several head replacements and causing several severe file crashes. Also during periods when one or more of the swapping devices was down, swapping off of moving head disks reduced efficiency substantially. These problems appear to have been solved since April. Other components of the system which have given trouble are the TYMNET interface (already mentioned), the PDP-10 memories and the moving head disks. The KI-10 CPU has been very stable and given only one problem over the past year (an I/O bus driver).

Most of the hardware problems have been very hard to track down as they caused crashes perhaps once per day and would not recur under diagnostic testing (in general TENEX exercises system components harder than do diagnostics). DEC has been very responsive in helping to find the problems in contrast to last year - the problems have simply gotten harder. It appears that the troubles should settle down soon as a number of intermittent faulty components have been found at last. We consider it a first order of business to improve these statistics.

From the user's viewpoint, besides the obvious inconvenience of not being able to work during down time, the fragility of the highly interlinked TENEX file system has caused several occasions of having to backup to previous file system states. We save changed files daily and copy the entire file system to fresh disk packs weekly. Thus an unexpected crash may cause the loss of up to one day's worth of work - it in fact may take longer for a given user to reconstruct the lost work if complex debugging or development changes were involved and undocumented. When the system is known to be subject to intermittent crashes, we backup more often to protect users. We are also investigating other modes of backup, now that we are on the ARPANET, such as the Datacomputer at the Computer Corporation of America.

Our current schedule for system backup is early Sunday morning (Pacific Time). We do not have enough staff for around the clock coverage and while we have overlapped staff to provide some weekend support and have scheduled backup then, this down-time for backup has been inconvenient for some users. We have tried to be responsive to these demands in that file backups used to be done to magnetic tape (requiring up to 6 hours for our size file system). We replaced this procedure with direct disk pack copying to reduce the time to about 2.5 hours. This eases the burden in down-time for essential backup operations; the next step would be to have enough staff to allow backup during very early morning hours to inconvenience (at least some) users less.

Another aspect of reliability and backup is the need to assure computing service for critical demonstrations, lectures, and the like. We are attempting to establish such a relationship with existing TENEX sites locally who are on the ARPANET but a surprisingly large number of problems arise; administrative and technical. These include the mechanics of moving files around beforehand (if a machine is down, files cannot be moved after the fact), allocation of space and time on otherwise heavily loaded machines, software compatibility in terms of monitor and languages, and approval of arrangements through responsible funding agencies. We are still working on this type of backup with an immediate need to support the AIM Workshop at Rutgers this June.

II.A.2.b USER SUPPORT AND INTERFACES

We have already addressed one aspect of user support from the system viewpoint; that of adapting system functions and defaults to individual users. The following are aspects of user support involving specific pieces of software made available and attempts to facilitate user contact with them.

Languages and Utility Programs:

A great deal of work was done during the past year in bringing up and improving the menu of subsystems available to users. New languages include SITBOL (Stevens Institute PDP-10 SNOBOL), FORTRAN-10 (release 4A), SAIL (TENEX version), TBASIC (Dartmouth language

definition with debugger, subroutines, etc.), INTERLISP updates, MACRO-10 update, FAIL update, ILISP (UC Irvine version of LISP 1.6), BCPL-10/11 (brought up by Rovner at Rochester), BLISS-10/11, and a preliminary version of PDP-11 SAIL (by Mr. Clark Wilcox of SUMEX). The PDP-11 SAIL compiler implements a significant subset of the SAIL language in a compiler which is highly machine independent by design. Code can be generated at present for the PDP-10, PDP-11, and IBM 360/370. Other machines are under consideration. The compiler itself currently runs on either the PDP-10 or a PDP-11/45 (with 32K of memory). Additional design information is given in Appendix D. In conjunction with these have come a variety of new utility programs including LINK-10, and CREF.

Beyond the language-related additions are a range of programs for text editing (including a CRT-oriented text editor, TV by Mr. Pentti Kanerva of IMSSS), mail handling, text justification (PUB and some new macro libraries), budgetting, typescript recording, multiple job fork control, mathematical modeling (MLAB from Gary Knott at NIH), graphics (OMNIGRAPH by Mr. R. Sproull at Xerox PARC), and so on. Rather than try to enumerate all of the available programs here, a brief summary of the major subsystems available is in Appendix E.

In a number of cases, considerable difficulties arose in importing software from various sites. These problems came about for a variety of reasons including getting incomplete sets of source files, programs written to take advantage of special system features and conventions not adopted universally, and inherent differences between TENEX and TOPS-10 (see "Compatibility" below, page 17).

Documentation and Education:

A substantial effort was made to better document subsystems and bugs for the programs available at SUMEX. As we have imported much of the software from other sources, we use available documentation where possible, update and adapt it where feasible, and write or rewrite from scratch where necessary. The reader is referred to Appendix D for a current listing of the <DOC> directory containing the on-line documents and a summary of available hard copy documents. Dr. Nancy Smith has completely revised the SOS and PUB documents and Dr. Robert Smith of IMSSS has prepared a document describing the TENEX version of SAIL. In addition numerous user help documents have been prepared for initial system access, network use, and subsystem usage aids.

Courses were prepared and given at Stanford covering the system assembly language (MACRO-10) and the TENEX system calls (JSYS's) (Messrs. Heathman and Crossland); TENEX SAIL (Dr. Robert Smith); and a class will be given shortly on PUB which is a powerful text justification language (Dr. N. Smith).

In addition to these more or less formal user support efforts, a large fraction of available staff time goes to tracking down real or perceived bugs encountered by users working on the system. It may be an under estimate to say that in excess of 30% of staff time is allocated to this purpose. Through the LINK and SNDMSG facilities, the staff provides help to users wherever they are located.

Compatibility Issues:

Over the past year, in our commitment to software importation where possible rather than reinvention, we have encountered the problems of software incompatibility between various machines and operating systems fairly often. This problem is present to some extent between various TENEX sites where different releases of the system or languages are run or where local system additions or changes create problems (some TOPS-10 systems or runtime programs are non-standard as well). It is felt most acutely, however, between TENEX and TOPS-10 systems. A number of the problems stem from operational issues; file names or locations may be "hardwired" into a program and the same convention does not apply at other sites or file search pathways and hierarchies are not identical at all sites. These problems may be quite baffling without source files or deep insight into the program. These kinds of problems point up examples of where improved programming practices would make exportation of software more easily managed.

More difficult problems arise over basic incompatibilities between the TENEX and TOPS-10 systems. These arise either because languages of the same name (ignoring modifiers such as "TENEX" and "DEC" which anticipate problems) are not really the same or because definitions and design features of the two systems are inherently different. Such problems have been particularly frustrating for some of the people and programs originating from the NIH-DCRT machines but may be present for any program designed to run on the TOPS-10 system.

One area of considerable effort by Dr. R. Smith over the past year has been in narrowing the differences between TOPS-10 SAIL and TENEX SAIL. He has implemented a number of default line editing options and pseudo-interrupt options (e.g., control O to terminate terminal output) so that these functions will be transparent between the two machines. Other areas are more difficult to deal with including file system structure and protection and system calls. The TENEX file system is more elaborate than the TOPS-10 system in a number of ways such as naming conventions and the accommodation of multiple versions of a file as well as procedures for getting rid of files. In other ways, such as protection, different conventions were adopted. Through proper choice of defaults within the TENEX directory specifications for a given user, the naming problem can be mitigated; however, long names will still be recognized by TENEX and not by TOPS-10. The protection differences cannot be completely fixed either as the mapping from one system to the other cannot be easily made in all cases.

The issue of system calls is another difficult area; in TENEX the system calls (JSYS's) are different than those in TOPS-10 (UUO's). They are implemented using the hardware, however, in non-interfering ways so that it is possible to write an emulator program (PA1050) in TENEX which traps all DEC-style calls and translates them into "equivalent" sequences of JSYS calls. The problems arise when a) there does not exist a functionally equivalent translation or b) DEC has created a new UUO (as continually happens with new TOPS-10 releases) which the emulator does not know about. Mr. J. Crossland

(and before him S. Reiss) has spent considerable effort in tracking down and repairing as consistently as possible these kinds of problems. A major difficulty is that the original PA1050 emulator was written for a much earlier DEC system and has grown and been updated in something of a "crazy quilt" fashion. We estimate 1 - 2 man-years of effort to properly redo the package.

Many of the compatibility problems can be avoided prospectively through proper programming practices. This does not alleviate the difficulties in adapting older programs - a conversion effort of some sort is necessary although once done, through conditional compilation statements (say in SAIL), future compatibility can be maintained while continuing development on only one copy of the source program. We will continue to work to minimize these headaches and remain available to advise and help users as much as possible. Despite these compatibility difficulties we feel that the choice of TENEX was the correct one for the AI mission of SUMEX-AIM, primarily because of the advantages of the demand paging LISP environment uniquely available in TENEX.

Library Building:

Another aspect of user community support and a key element in the community-oriented mandate of SUMEX-AIM is the assimilation of software tools from active groups within and without the immediate SUMEX user groups. We have begun an effort to accumulate useful SAIL library routines from the various groups which have been working with this language (Stanford AI, IMSSS, SRI, NIH, USC-ISI, etc.). It is somewhat surprising that so little communication of SAIL library programs has taken place - it is almost literally true that each user has his own stock of tools in private procedure libraries. We have sent a letter to interested groups soliciting inputs on a basis which attempts to balance the problem of assuring library quality and integrity against establishing so high a threshold for quality and polish that individuals are not motivated to cooperate. This effort has just recently begun and no results are reportable to date.

II.A.3 RESOURCE MANAGEMENT

Over the past year, the SUMEX project has devoted a substantial part of its effort toward its community-building role in recruiting new project, promoting interactions between user projects, and encouraging dissemination of running performance programs to medical scientists. A representative summary of SUMEX's community orientation in outlook and in action is given in a paper on networking and collaborative research (see Appendix F). This paper will be presented at the 170th American Chemical Society symposium this August 1975 and will appear in the proceedings.

The following summarizes specific aspects of SUMEX-AIM community management activities.

Dedication:

The SUMEX resource, having reached fully operational status by early fall 1974, held a dedication program at Stanford University on November 14, 1974. The program was an all-day symposium with the morning devoted to technical presentations by the initially authorized projects (see Section IV: DENDRAL, RUTGERS, MYCIN, Higher Mental Functions Modeling, and Protein Structure Modeling). The afternoon session addressed more global policy issues related to resource sharing and included presentations by Dr. Lederberg (SUMEX Principal Investigator), Dr. Thomas Bowery (Director of the NIH Division of Research Resources), Dr. W. Miller (Provost of Stanford University), and Dr. J.G.R. Licklider (Director of ARPA's Information Processing Techniques Office).

In addition to the program at Stanford, the attendant press releases and handout brochure (Appendix H), we also published an announcement of the SUMEX resource in the September 1974 SIGART (Special Interest Group for Artificial Intelligence) newsletter of the Association of Computing Machinery (ACM) and made a series of presentations on SUMEX and its related projects at the SIGBIO session of the 1974 annual ACM conference in San Diego (November 13, 1974).

Management Committees:

The SUMEX-AIM resource is constituted to attempt to bring into closer contact collaborating health research groups from around the country. This mission entails both the recruitment of appropriate research projects interested in medical AI applications and the catalysis of interactions among these groups and the broader medical community. As this effort is not a unilateral undertaking by its very nature, we have created several management committees to assist in administering the various portions of the SUMEX resource. As defined in the SUMEX-AIM management plan adopted at the time the resource grant was awarded, the available facility capacity is allocated 40% to Stanford Medical School projects, 40% to national projects, and 20% to system development and related functions.

Within the Stanford aliquot, Dr. Lederberg has established an advisory committee to assist him in selecting and allocating resources among projects appropriate to the SUMEX mission. The current membership of this committee is listed in Appendix G.

For the national community, two committees serve complementary functions. An Executive Committee oversees the operations of the resource as related to national users and makes the final decisions on authorizing admission for projects. It also establishes policies for resource allocation and approves plans for resource development and augmentation within the national portion of SUMEX. The Executive Committee oversees the planning and implementation of the AIM Workshop series and assures coordination with other AIM activities as well. The workshops are being carried out under Dr. S. Amarel of the Rutgers Computers in Biomedicine resource. The current membership of the Executive committee is listed in Appendix G.

Under the Executive Committee functions an Advisory Group representing contact with medical and computer science research relevant to AIM goals. The Advisory Group serves several functions in advising the Executive Committee; 1) recruiting appropriate medical/computer science projects, 2) reviewing and recommending priorities for allocation of resource capacity to specific projects based on scientific quality and medical relevance, and 3) recommending policies and development goals for the resource. The current Advisory Group membership is given in Appendix G.

These committees are actively functioning in support of the resource. Meetings to date have been held by telephone conference for the most part owing to the size of the groups and the difficulties in arranging for travel to meet face to face. These "missings" (a term coined by Dr. Licklider), in conjunction with terminal access to related text materials, have served quite well in accomplishing the agenda business and facilitate greatly the arrangement of meetings. A few technical problems occasionally attend such sessions such as poor telephone reception for some members but in general this approach is quite satisfactory.

New Project Recruiting:

As a result of the public announcements of the SUMEX resource, NIH reviews of the Health Manpower Act (769-A) proposals, and personal contacts by the staff or committee members, a number of additional projects have been admitted to SUMEX; others are working tentatively as pilot projects or are under review. We have prepared a variety of materials for the new user ranging from general information such as is contained in the brochure (Appendix H) to more detailed information and guidelines for determining whether a user project is appropriate for the SUMEX-AIM resource. Dr. E. Levinthal has prepared a questionnaire to assist users seriously considering applying for access to SUMEX-AIM (see Appendix I). Pilot project categories have been established both within the Stanford and national aliquots of the facility capacity to assist and encourage projects just formulating possible AIM proposals pending a formal review.

The projects newly admitted over the past year include (see Section IV for more detailed descriptions):

Stanford - (Pilot)

- 1) Information Processing Psychology; Drs. E. Feigenbaum (Stanford) and H. Cohen (UC San Diego)

National -

- 1) Diagnostic Logic Project (DIALOG); Dr. H. Pople and J. Myers, M.D. (University of Pittsburgh)
- 2) Medical Information Systems Lab (MISL); Dr. B. McCormick and M. Goldberg, M.D. (University of Illinois at Chicago Circle)

- 3) Distributed Data Base System for Chronic Diseases; Drs. F. Kuo (University of Hawaii), R. Nordyke, M.D. (Pacific Health Research Institute), and Dr. C. Kulikowski (Rutgers University)

We believe, within the current system capacity and with proper scheduling and capacity allocation controls, that another 2 or 3 major projects can be accommodated plus 5 or 6 minor projects. Here major and minor magnitude refers only to amount of computer resource consumption and not to scientific quality. Clearly the admission of these additional projects is based upon the ability to direct system use to currently underloaded parts of the day. This may require management committee decisions making access for some projects conditional upon system use at non-peak periods as well as other measures to encourage load leveling.

For another perspective on the community of projects currently being supported by the resource, see Appendix J. This appendix contains material prepared in response to a congressional inquiry to NIH-BRB on the scope and cost of community support by the SUMEX resource.

As an additional aid to new projects or collaborators with existing projects, we have a limited amount of funds which are being used to support terminals and communications needs of users without access to such equipment. We are currently leasing 5 terminals and 3 modems for users and will be providing some foreign exchange lines to users to improve network response time.

Utility of Intergroup Coupling:

One of the central objectives of the SUMEX resource is to encourage routine contact between remote groups. This may manifest itself in a number of ways such as collaboration within a project between researchers who are not geographically close, interactions between research projects which are at separate institutions, and dissemination of research products to users not close to the necessary specialized facilities. We are developing examples of useful collaboration in all of these categories as is summarized in the individual project descriptions attached in Section IV.

Several of the approved projects already involve remote collaborations; Rutgers Computers in Biomedicine (between Rutgers University, Mt. Sinai Hospital in New York, Johns Hopkins University in Baltimore, and Washington University in St. Louis), Protein Structure Modeling (between Stanford University and UC San Diego), and Distributed Data Bases (between the University of Hawaii and Rutgers University). The following message quoted from the Protein Structure Modeling group points up the utility of network relationships for coordinating remote development activities:

Date: 2 JAN 1975 0010-PST

From: ENGELMORE

Subject: ADVANTAGES OF SUMEX FOR COLLABORATIVE RESEARCH

"Yesterday, I was engaged for several hours in a very interesting collaboration, involving SUMEX, Steve Freer and others at UCSD. I was debugging a program which Steve recently sent to me, and running into a variety of bewilderments. We then linked to each other, so that my program output came out on Steve's terminal as well as mine. He then would comment on the output, direct my attention to appropriate parts of the program, and suggest changes. We made remarkable progress in that mode; it was as efficient as having Steve and some of his colleagues sitting right next to me as I worked. Although I knew perfectly well that networks and links permit this mode of operation, actually doing it was a fascinating experience. For Freer, however, it was a revelation! He had no idea before this that two people, 500 miles apart, could both examine program output independently and simultaneously. It really turned him on.

Had I not been able to converse with the UCSD group in "real time", I very likely would have traveled to La Jolla and worked there. So I feel the system we have in SUMEX is a real time and energy saver."

In the second category we are also developing examples of mutually useful interactions between research groups. Because the programs are accessible through common communication services, remote interactive criticism and discussions are possible as the programs are being developed. The following note describing an interaction between the MYCIN group at Stanford and the DIALOG group at Pittsburgh illustrates the point:

Date: 14 MAR 1975 1903-PST
 From: SHORTLIFFE
 Subject: Demo Last Saturday

"Bruce Buchanan suggested that I tell you about a use of the SUMEX system that we experimented with last Saturday. Harry Pople's group at Pittsburgh was interested in getting some reaction to their DIALOG system, so we arranged a time last Saturday morning for a demonstration. Meanwhile, several members of the medical diagnostic group at Rutgers were also interested and asked to sit in. We therefore all linked to one another at a prearranged time, and for about 2 or 3 hours, Pople demonstrated their program and then watched while I ran it on a patient of my own choosing. A number of comments and questions arose which were easily handled by the link procedure, and when the demonstration was over we continued to discuss via the link a number of other topics of mutual interest including plans for the AIM conference at Rutgers in June. It was a very satisfactory way to 'meet' without burning up the long distance phone lines (they, of course, were all logged on via the TYMNET), and the incident

may therefore be of interest to you when you discuss some of the novel advantages of a national resource such as SUMEX."

Another form of intergroup collaboration is developing between the Rutgers project and the MISL project at Illinois. The Illinois group is planning to use the Rutgers glaucoma programs as an integral part of their research with the University of Illinois Eye Clinic. There have been a number of delays in getting this interaction working smoothly caused by the problems in getting network connections working, needed language support debugged, and finally getting the glaucoma programs to a state where routine access is possible for the Illinois project. These should be solved now and hopefully the next report will see an active collaboration between these groups.

Finally, those projects with programs beyond the early development stages (principally DENDRAL, MYCIN, Rutgers glaucoma, and Higher Mental Functions PARRY) have made substantial investments in liaison and programmer time to facilitate non-expert user interfaces to their performance programs. These resulting programs have then been made available to selected professionals outside of the development groups for experimental use and appraisal. In numerous cases, the network connections have allowed contacts with these users from areas quite remote from Stanford and where it would be impossible to mount the programs for lack of necessary specialized computing facilities. These contacts have produced promising results even at these early stages as described in the individual project summaries (Section IV). A major objective of the SUMEX project community is to continue establishing contacts with non-computer scientists in the various research areas under investigation and to demonstrate and evaluate the utility of the medical AI programs.

Resource Allocation Policies:

As the SUMEX facility becomes increasingly loaded, a number of diverse and conflicting demands can be identified which require controlled allocation of critical facility resources (file space and central processor time). We have already spelled out a policy for file space management; an allocation of file storage will be defined for each authorized project in conjunction with the management committees. This allocation can be divided among project members in any way desired by the individual principal investigators. No system allocation enforcement will be imposed as long as there is adequate file capacity left in the system to afford as much flexibility as possible to projects for temporary file space needs. However, when used space approaches system capacity, a variety of tools (verbal requests, deleted file expunging, and forced file archival) are available to ensure that projects observe their allocated space. So far the user community has been very cooperative and has responded to verbal requests for file space clean-up.

As described under "System Development Progress", we have

implemented a primitive CPU scheduling algorithm intended to ensure that no one user gets more than a fair share of the machine when other users are contending. As discussed there, it is likely that a more sophisticated scheme will be necessary to meet the community needs (fixed personal schedules and relative priority ratings). This may be implemented by some form of "reservation" system where some prescribed fraction of the machine can be expected for a given individual or project during a specified period at the expense of priority at other times. We are discussing these issues with the management committees to evolve the most beneficial policy for the SUMEX-AIM community.

As also mentioned earlier, we are developing a categorization of users in terms of access privileges. These range from fully authorized users to guests and network visitors in descending order of system capabilities. We want to encourage bona fide medical people to experiment with the various programs available with a minimum of red tape while not allowing unauthenticated users to bypass the advisory group screening procedures by coming on as guests. We will continue developing this mechanism in conjunction with management committee policy decisions.

AIM Workshop Support:

The Rutgers Computers in Biomedicine resource (under Dr. Saul Amarel) is actively working on plans for the first AIM workshop this June. The current plans call for a one day general session covering a range of topics related to artificial intelligence research, medical needs, and resource sharing policies within NIH. The following three days will include a more intimate set of working sessions to allow first hand experience with running programs for various prospective users and interested research people. The SUMEX facility will act as the computing base for the workshop demonstrations. We are in the process of working with Rutgers to provide backup modes for program demonstrations in the event of system failure.

II.A.4 FUTURE PLANS

System Performance:

In the next year we will work on improving system performance based on measurement data now being collected and evaluated. For example, we want to tune the working set size limits and logic to improve the trade-off between paging traffic and the number of jobs in core. We are working on implementing an algorithm to more efficiently utilize swapping storage by migrating dormant pages off to moving head storage. In parallel with our measurement efforts, other groups (USC-ISI) are debugging and testing TENEX systems with memories larger than 256K words.

These efforts will assist us to plan where key augmentations (memory, CPU, swapping storage, file access) could increase throughput as the AIM community grows. We are currently developing a plan to

overcome the CPU bottleneck which we feel will shortly become the most critically limited resource. Preliminary details of this plan are described under "Equipment" in next year's budget (see page 51). We will refine this plan and submit it to the Executive Committee and BRB for approval. We have requested that any unspent funds from year 02 be carried forward to year 03 for this purpose.

We will investigate bringing up version 1.33 of TENEX with necessary KI-10 modifications in order to stay current with other TENEX sites (and to facilitate maintaining an up-to-date INTERLISP subsystem) as well as evaluate the scheduler and resource allocation group features introduced by BB&N in 1.33 to see if they will assist the allocation controls contemplated for various user needs. To date KI-TENEX 1.33 is still being debugged at NASA-AMES and we will wait until the system runs smoothly and reliably before experimenting with it.

We plan to bring up a batch processing capability for those jobs which need not run interactively. A primitive system has been put together by USC-ISI and will be extended where needed (e.g., to allow multiple jobs, priority control so as not to compete with interactive work, etc.). In addition, we will add a hardcopy plotter (plotter available from another project) to the system along with a spooler to facilitate multiple use.

We will continue to refine the Executive program and capabilities for guest users.

We will also investigate ways of improving network communication services. This will include attempts to optimize our current facilities for users through better ties to the networks and selective lines to tie individual users into more advantageous access points. We will also continue to explore other network and communication alternatives as they become available over the next year. Specific goals include improved response times and increased output speeds. We expect the ARPANET link to improve about mid-summer with the addition of the other 50 K baud line to the Stanford AI Laboratory IMP.

Adaptive User Interfaces:

We plan to continue work toward a more adaptive system for users including both simplifying access for non-expert users and anticipating default parameter conventions of individual users. Longer term planning may look at more sophisticated user modeling and the possibility of putting such personalized interfaces into a user's terminal. We are now in the process of defining system calls which will make user information uniformly accessible to programs that choose to make use of it.

Software Facilities and Libraries:

There is a continuing need for improved documentation of various aspects of the system and of available programs. We will be up-